



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/808,561

03/24/2004

Ofir Zohar

ASSIA 21.061

4776

26304

7590

09/21/2006

KATTEN MUCHIN ROSENMAN LLP
575 MADISON AVENUE
NEW YORK, NY 10022-2585

EXAMINER

GOLDEN, JAMES R

ART UNIT

PAPER NUMBER

2187

DATE MAILED: 09/21/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No. 10/808,561	Applicant(s) ZOHAR ET AL.	
	Examiner James Golden	Art Unit 2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 24 March 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-33 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-33 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 24 March 2004 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| Paper No(s)/Mail Date <u>24 March 2004</u> | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

The instant application 10/808561 has a total of 33 claims pending. There are 6 independent claims and 27 dependent claims.

Information Disclosure Statement

1. The information disclosure statement submitted on 03/24/2004 is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Priority

2. Applicant's claim for the benefit of a prior-filed application under 35 U.S.C. 119(e) or under 35 U.S.C. 120, 121, or 365(c) is acknowledged, and priority is granted to the filing date 07/15/2003 of parent applications 10/620080 and 10/620249.

Claim Rejections - 35 USC § 112

3. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

4. **Claims 27, 29, 31 and 33** are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

The term "fast-access-time" in claims 27, 29, 31 and 33 is a relative term which renders the claim indefinite. The term "fast-access-time" is not defined by the claim, the specification does not provide a standard for ascertaining the requisite degree, and one

Art Unit: 2187

of ordinary skill in the art would not be reasonably apprised of the scope of the invention. This rejection could be overcome by specifying a particular speed of access or by stating a particular fast-access-time mass storage device.

Claim Rejections - 35 USC § 102

5. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

6. **Claims 1-8, 12-21 and 25** are rejected under 35 U.S.C. 102(e) as being anticipated by Hicken et al. (US 2004/0153727).

7. **With respect to claim 1**, Hicken et al. disclose a method for managing a data storage system (**300 of Fig. 3; paragraph 0038, lines 4-10**), comprising:

- configuring a first cache (**339 of Fig. 3**) to perform at least one of the operations of retrieving data from and storing data at a first range of logical addresses (LAs) in a storage device (**paragraph 0038, lines 13-17; paragraph 0039, lines 7-10; paragraph 0041, lines 10-14; primary cache 333 of storage controller 370-1 is the cache for LA1, and secondary cache 339 is the redundant cache for LA1, and resides on storage controller 370-2**);

Art Unit: 2187

- configuring a second cache (**333 of Fig. 3**) to perform at least one of the operations of retrieving data from and storing data at the first range of LAs (**paragraph 0039, lines 5-7; paragraph 0041, lines 10-14**);
- configuring one or more third caches (**338 of Fig. 3**) to perform at least one of the operations of retrieving data from and storing data at a second range of LAs in the storage device (**paragraph 0041, lines 14-17**);
- detecting an inability of the second cache to retrieve data from or store data at the first range of LAs (**paragraph 0042, lines 7-9; when the storage controller 370-1 fails, cache memory 339 fails as well**); and
- reconfiguring at least one of the first cache and the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs (**paragraph 0042, lines 9-13; cache 339 becomes the cache for LA1**).

8. **With respect to claim 2**, Hicken et al. disclose the method according to claim 1 (see above paragraph 7), and comprising configuring one or more interfaces (**CPUs 331 and 336 of Fig. 3**) to receive input/output (IO) requests (**paragraph 0025, lines 4-6 [a distinct but almost identical embodiment]; paragraph 0039, lines 15-16**) from host processors (**310 of Fig. 3; paragraph 0038, lines 4-7**) directed to specified LAs (**paragraph 0025, lines 3-6**) and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs (**paragraph 0039, lines 15-16**).

Art Unit: 2187

9. **With respect to claim 3**, Hicken et al. disclose the method according to claim 2 (see above paragraph 8), wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs (**paragraph 0038, lines 13-17**), and wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping (**paragraph 0025, lines 4-6; paragraph 0026, lines 1-2; the CPUs onboard the storage controllers receive storage requests from the host, and issue them to the caches**), and wherein detecting the inability comprises generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs (**paragraph 0042, lines 9-13**), and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping (**paragraph 0042, lines 9-13**).
10. **With respect to claim 4**, Hicken et al. disclose the method according to claim 1 (see above paragraph 7), wherein reconfiguring the at least one of the first cache and the one or more third caches comprises processing data in the first cache and the one or more third caches so as to restore the first cache and the one or more third caches to a state of full data redundancy (**paragraph 0044, lines 15-24**).
11. **With respect to claim 5**, Hicken et al. disclose the method according to claim 4 (see above paragraph 10), wherein processing the data comprises classifying data in the first cache into a plurality of data groups (**paragraph 0043, lines 10-15; when the caches are flushed, only the dirty data is flushed, not the resident data, which shows the data was classified**).

12. **With respect to claim 6**, Hicken et al. disclose the method according to claim 5 (see above paragraph 11), wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the one or more third caches (**paragraph 0043, lines 10-15; dirty data is stored on the third cache 338**).
13. **With respect to claim 7**, Hicken et al. disclose the method according to claim 5 (see above paragraph 11), wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the storage device (**paragraph 0043, lines 10-15; dirty data is flushed to the storage units**).
14. **With respect to claim 8**, Hicken et al. disclose the method according to claim 1 (see above paragraph 7), wherein reconfiguring the at least one of the first cache and the one or more third caches comprises retaining an initial configuration of the first cache (**paragraph 0042, lines 9-13, lines 18-22; after the second cache 333 fails, the first cache 339 retains its initial configuration and is now used to address storage requests for LA1**).
15. **With respect to claim 12**, Hicken et al. disclose the method according to claim 1 (see above paragraph 7), and comprising providing a system manager (**host computer 310 and CPUs 331 and 336 of Fig. 3**) which is adapted to configure the first, second and one or more third caches (**paragraph 0025, lines 4-6; paragraph 0039, lines 15-16**), to detect the inability (**paragraph 0028, lines 1-4**), and to reconfigure the at least one of the first cache and the one or more third caches (**paragraph 0042, lines 9-13, lines 18-22**).
16. **With respect to claim 13**, Hicken et al. disclose the method according to claim 12 (see above paragraph 15), wherein providing the system manager comprises

Art Unit: 2187

incorporating one or more manager processing units into at least one of the storage device, the first cache, the second cache, and the one or more third caches (**host computer 310 and CPUs 331 and 336 of Fig. 3**), and operating the one or more manager processing units in a cooperative manner (**paragraph 0040; all of the CPUs are connected and work together**).

17. **With respect to claim 14**, Hicken et al. disclose a data storage system, comprising:

- a storage device (**300 of Fig. 3; paragraph 0038, lines 4-10**) wherein data is stored at logical addresses (LAs);
- a first cache (**339 of Fig. 3**) which is configured to perform at least one of the operations of retrieving data from and storing data at a first range of LAs in the storage device (**paragraph 0038, lines 13-17; paragraph 0039, lines 7-10; paragraph 0041, lines 10-14; primary cache 333 of storage controller 370-1 is the cache for LA1, and secondary cache 339 is the redundant cache for LA1, and resides on storage controller 370-2**);
- a second cache (**333 of Fig. 3**) which is configured to perform at least one of the operations of retrieving data from and storing data at the first range of LAs (**paragraph 0039, lines 5-7; paragraph 0041, lines 10-14**);
- one or more third caches (**338 of Fig. 3**) which are configured to perform at least one of the operations of retrieving data from and storing data at a second range of LAs in the storage device (**paragraph 0041, lines 14-17**); and
- a system manager (**host computer 310 and CPUs 331 and 336 of Fig. 3**)

- which is adapted to detect an inability of the second cache to retrieve data from or store data at the first range of LAs (**paragraph 0042, lines 7-9; when the storage controller 370-1 fails, cache memory 339 fails as well**), and
- which reconfigures at least one of the first cache and the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs (**paragraph 0042, lines 9-13; cache 339 becomes the cache for LA1**).

18. **With respect to claim 15**, Hicken et al. disclose the storage system according to claim 14 (see above paragraph 17), and comprising one or more interfaces (**CPUs 331 and 336 of Fig. 3**) which are configured to receive input/output (IO) requests (**paragraph 0025, lines 4-6 [a distinct but almost identical embodiment]; paragraph 0039, lines 15-16**) from host processors (**310 of Fig. 3; paragraph 0038, lines 4-7**) directed to specified LAs (**paragraph 0025, lines 3-6**) and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs (**paragraph 0039, lines 15-16**).

19. **With respect to claim 16**, Hicken et al. disclose the storage system according to claim 15 (see above paragraph 18), wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs (**paragraph 0038, lines 13-17**), and wherein the one

Art Unit: 2187

or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping (**paragraph 0025, lines 4-6; paragraph 0026, lines 1-2; the CPUs onboard the storage controllers receive storage requests from the host, and issue them to the caches**), and wherein detecting the inability comprises the system manager generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs (**paragraph 0042, lines 9-13**), and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping (**paragraph 0042, lines 9-13**).

20. **With respect to claim 17**, Hicken et al. disclose the storage system according to claim 14 (see above paragraph 17), wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the first cache processing data therein and the one or more third caches processing data therein so as to restore the first cache and the one or more third caches to a state of full data redundancy (**paragraph 0044, lines 15-24**).

21. **With respect to claim 18**, Hicken et al. disclose the storage system according to claim 17 (see above paragraph 20), wherein processing the data comprises classifying data in the first cache into a plurality of data groups (**paragraph 0043, lines 10-15; when the caches are flushed, only the dirty data is flushed, not the resident data, which shows the data was classified**).

22. **With respect to claim 19**, Hicken et al. disclose the storage system according to claim 18 (see above paragraph 11), wherein one of the data groups comprises dirty

Art Unit: 2187

data, and wherein processing the data comprises storing the dirty data at the one or more third caches (**paragraph 0043, lines 10-15; dirty data is stored on the third cache 338**).

23. **With respect to claim 20**, Hicken et al. disclose the storage system according to claim 5 (see above paragraph 11), wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the storage device (**paragraph 0043, lines 10-15; dirty data is flushed to the storage units**).

24. **With respect to claim 21**, Hicken et al. disclose the storage system according to claim 1 (see above paragraph 7), wherein reconfiguring the at least one of the first cache and the one or more third caches comprises retaining an initial configuration of the first cache (**paragraph 0042, lines 9-13, lines 18-22; after the second cache 333 fails, the first cache 339 retains its initial configuration and is now used to address storage requests for LA1**).

25. **With respect to claim 25**, Hicken et al. disclose the storage system according to claim 14 (see above paragraph 17), wherein the system manager comprises one or more manager processing units incorporating one or more manager processing units into at least one of the storage device, the first cache, the second cache, and the one or more third caches (**host computer 310 and CPUs 331 and 336 of Fig. 3**), and wherein the one or more manager processing units operate in a cooperative manner (**paragraph 0040; all of the CPUs are connected and work together**).

26. **Claims 26-33** are rejected under 35 U.S.C. 102(e) as being anticipated by Henry et al. (US 6,898,666).

27. **With respect to claim 26**, Henry et al. teach a storage system, comprising:

Art Unit: 2187

- one or more mass storage devices, coupled to store data at respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23, lines 55-67**);
- a plurality of interim fast-access-time caches (**cache pools 1 and 2**), configured to operate independently of one another, each interim fast-access-time device being assigned a respective second range of the LAs (**column 5, lines 45-55**); and coupled to receive data from and provide data to the one or more -mass-storage devices having LAs within the respective second range (**column 4, lines 55-59**); and
- one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the requests to the interim fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).

28. **With respect to claim 27**, Henry et al. teach the storage system according to claim 27 (see above paragraph 27), wherein the mass storage devices comprise fast-access-time mass storage devices (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

29. **With respect to claim 28**, Henry et al. teach a method for storing data, comprising:

- storing the data in one or mass storage devices having respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23 & 55-67**);
- assigning to each of a plurality of interim-fast-access-time nodes (**cache pools 1 and 2**), configured to operate independently of one another, a respective second range of the LAs (**each cache pool is assigned separate LBA ranges for the associated disks; column 5, lines 45-55**);

Art Unit: 2187

- coupling the plurality of interim-fast-access-time caches to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range (**column 4, lines 55-59**);
- receiving input/output (IO) requests from host processors directed to specified LAs (**column 2, lines 11-15; column 4, lines 64-67**); and
- directing all the IO requests to the interim-fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).

30. **With respect to claim 29**, Henry et al. teach the method according to claim 28 (see above paragraph 29), wherein the mass storage devices comprise fast-access-time mass storage devices (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

31. **With respect to claim 30**, Henry et al. teach a system for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23, lines 55-67**), comprising:

- a plurality of interim fast-access-time caches (**cache pools 1 and 2**), configured to operate independently of one another, each interim fast-access-time device being assigned a respective second range of the LAs (**column 5, lines 45-55**); and coupled to receive data from and provide data to the one or more -mass-storage devices within the respective second range (**column 4, lines 55-59**); and
- one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the requests to the interim fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).

Art Unit: 2187

32. **With respect to claim 31**, Henry et al. teach the storage system according to claim 30 (see above paragraph 29), wherein the mass storage devices comprise fast-access-time mass storage devices (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

33. **With respect to claim 32**, Henry et al. teach a method for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23, lines 55-67**), comprising:

- assigning to each of a plurality of interim-fast-access-time caches (**cache pools 1 and 2**), configured to operate independently of one another, a respective second range of the LAs (**each cache pool is assigned separate LBA ranges for the associated disks; column 5, lines 45-55**);
- coupling the plurality of interim-fast-access-time caches to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range (**column 4, lines 55-59**);
- receiving input/output (IO) requests from host processors directed to specified LAs (**column 2, lines 11-15; column 4, lines 64-67**); and
- directing all the IO requests to the interim-fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).

34. **With respect to claim 33**, Henry et al. teach the method according to claim 28 (see above paragraph 29), wherein the mass storage devices comprise fast-access-time mass storage devices (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

Claim Rejections - 35 USC § 103

35. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

36. **Claims 9-11 and 22-24** are rejected under 35 U.S.C. 103(a) as being unpatentable over Hicken et al. (US 2004/0153727) in view of Karger et al. ("Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web," by in the Proceedings of the 29th ACM Symposium on Theory of Computing, Pages 654-663).

37. **With respect to claim 9**, Hicken et al. disclose the method according to claim 1 (see above paragraph 7). Hicken et al. do not disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

However, Karger et al. disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches (**page 5, section 4, "Consistent Hashing"**).

Hicken et al. and Karger et al. are analogous art because they are from the same field of endeavor, namely data caching.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the minimum redistribution in the form of consistent hashing of Karger et al. with the data caching redundancy system of Hicken et al. The motivation for doing so would have been because to prevent requiring a central server to distribute a completely updated hash table to all the machines every time a new machine is added to the network (**page 2, column 2, paragraph 2 beginning with “Our second...”**).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Karger et al. with Hicken et al. for the benefit of a data caching system with consistent hashing to obtain the invention as specified in claim 9.

38. **With respect to claim 10**, Hicken et al. disclose the method according to claim 9 (see above paragraph 37). Hicken et al. do not disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

However, Karger et al. disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function (**page 5, section 4, “Consistent Hashing”**).

39. **With respect to claim 11**, Hicken et al. disclose the method according to claim 9 (see above paragraph 37). Hicken et al. do not disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function.

However, Karger et al. disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function (**page 2, column 1, paragraph 5, “Our first tool, *random cache trees*...”**).

Art Unit: 2187

40. **With respect to claim 22**, Hicken et al. disclose the storage system according to claim 14 (see above paragraph 17). Hicken et al. do not disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

However, Karger et al. disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches (**page 5, section 4, “Consistent Hashing”**).

41. **With respect to claim 23**, Hicken et al. disclose the storage system according to claim 22 (see above paragraph 40). Hicken et al. do not disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

However, Karger et al. disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function (**page 5, section 4, “Consistent Hashing”**).

42. **With respect to claim 24**, Hicken et al. disclose the storage system according to claim 22 (see above paragraph 40). Hicken et al. do not disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function.

Art Unit: 2187

However, Karger et al. disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function (**page 2, column 1, paragraph 5, "Our first tool, *random cache trees*..."**).

Conclusion

43. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

- Chatterjee et al. (US 2005/0182906) teach a RAID system with redundant cache memories in RAID controllers;
- Ash et al. (US 2004/0059870) teach a pair of disk controller with redundant caches; and
- Hicken et al. (US 2003/0212864) teach a a pair of disk controller with redundant caches.

44. Any inquiry concerning this communication or earlier communications from the examiner should be directed to James Golden whose telephone number is 571-272-5628. The examiner can normally be reached on Monday-Friday, 8:30 AM - 5:30 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on 571-272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.


Art Unit: 2187

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

James R. Golden
Patent Examiner
Art Unit 2187



September 15, 2006



DONALD SPARKS
SUPERVISORY PATENT EXAMINER